



US008819690B2

(12) **United States Patent**
Alexander et al.

(10) **Patent No.:** **US 8,819,690 B2**

(45) **Date of Patent:** ***Aug. 26, 2014**

(54) **SYSTEM FOR REDUCING DATA TRANSFER LATENCY TO A GLOBAL QUEUE BY GENERATING BIT MASK TO IDENTIFY SELECTED PROCESSING NODES/UNITS IN MULTI-NODE DATA PROCESSING SYSTEM**

(75) Inventors: **Benjamin G. Alexander**, Austin, TX (US); **Gregory H. Bellows**, Austin, TX (US); **Joaquin Madruga**, Austin, TX (US); **Barry L. Minor**, Austin, TX (US)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1274 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **12/649,667**

(22) Filed: **Dec. 30, 2009**

(65) **Prior Publication Data**

US 2011/0161975 A1 Jun. 30, 2011

(51) **Int. Cl.**
G06F 9/46 (2006.01)

(52) **U.S. Cl.**
USPC **718/104**

(58) **Field of Classification Search**
USPC 718/104
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,137,115 B2 * 11/2006 Sakamoto et al. 718/102
7,155,716 B2 * 12/2006 Hooman et al. 718/102
7,877,751 B2 * 1/2011 Maeda et al. 718/102
2011/0161974 A1 * 6/2011 Kurabayashi et al. 718/104

FOREIGN PATENT DOCUMENTS

JP 02242434 * 9/1999 G06F 9/46

OTHER PUBLICATIONS

Tadashi Kamiwaki et al., Task Scheduling Method, Sep. 26, 1999.*

* cited by examiner

Primary Examiner — Thomas Lee

Assistant Examiner — Xuxing Chen


(74) *Attorney, Agent, or Firm* — Yudell Isidore Ng Russell PLLC

(57) **ABSTRACT**

A system for efficient dispatch/completion of a work element within a multi-node data processing system. The system comprises a processor performing the functions of: selecting specific processing units from among the processing nodes to complete execution of a work element that has multiple individual work items that may be independently executed by different ones of the processing units; generating an allocated processor unit (APU) bit mask that identifies at least one of the processing units that has been selected; placing the work element in a first entry of a global command queue (GCQ); associating the APU mask with the work element in the GCQ; and responsive to receipt at the GCQ of work requests from each of the multiple processing nodes or the processing units, enabling only the selected specific ones of the processing nodes or the processing units to be able to retrieve work from the work element in the GCQ.

19 Claims, 9 Drawing Sheets

GLBL CMD QUEUE WORK ITEM DISPATCH 400






FIG. 1





FIG. 2
PRIOR ART

FIG. 3



GLOBAL QUEUE WORK ITEM DISPATCH 400




FIG. 4

GCQ
NODE_0 APPROVED PROCESSORS 464A C_0 1000000000 C_3 0001000000

GCQ
NODE_1 APPROVED PROCESSORS 464B C_1 0000001000 C_3 0000100000

=

=

APU MASK 33D 1001010100

APU MASK 33D 1001010100

FIG. 5A

FIG. 5B

&

&

NODE_0 (PU) MASK 462A C_0 1000000000 C_1 0100000000 C_2 0010000000 C_3 0001000000

NODE_1 (PU) MASK 462B C_0 0000100000 C_1 0000010000 C_2 0000001000 C_3 0000000100

GLOBAL/Common CMD QUEUE
320

ND RANGE WORK ELEMENTS 325	SEEN CNTR 627	WIR CNTR 629	MASK BITS 330	ATTR 631
WE_R	NCU	NWI	nnnNNN	AI..AN
⋮	⋮	⋮	⋮	⋮
WE_2	48	2000	0011010111	AI..AN
WE_1	7	100	1011010111	AI..AN
WE_0	4	16	1001010100	AI..AN

FIG. 6A

GLOBAL/Common CMD QUEUE
320

ND RANGE WORK ELEMENTS 325	SEEN CNTR 627	WIR CNTR 629	MASK BITS 330	ATTR 331
WE_R	NCU	NWI	nnnNNN	AI..AN
⋮	⋮	⋮	⋮	⋮
WE_2	48	2000	0011010111	AI..AN
WE_1	7	100	1011010111	AI..AN
WE_0	4	12	1001010100	AI..AN

FIG. 6B

GLOBAL/Common CMD QUEUE
320

ND RANGE WORK ELEMENTS 325	SEEN CNTR 627	WIR CNTR 629	MASK BITS 330	ATTR 631
WE_R	NCU	NWI	nnnNNN	AI..AN
⋮	⋮	⋮	⋮	⋮
WE_2	48	2000	0011010111	AI..AN
WE_1	7	100	1011010111	AI..AN
WE_0	3	0	1001010100	AI..AN

FIG. 6C

GLOBAL/Common CMD QUEUE
320

ND RANGE WORK ELEMENTS 325	SEEN CNTR 627	WIR CNTR 629	MASK BITS 330	ATTR 631
WE_R	NCU	NWI	nnnNNN	AI..AN
⋮	⋮	⋮	⋮	⋮
WE_2	48	2000	0011010111	AI..AN
WE_1	7	90	1011010111	AI..AN
WE_0	1	0	1001010100	AI..AN

FIG. 6D




FIG. 7




Fig. 8

1

**SYSTEM FOR REDUCING DATA TRANSFER
LATENCY TO A GLOBAL QUEUE BY
GENERATING BIT MASK TO IDENTIFY
SELECTED PROCESSING NODES/UNITS IN
MULTI-NODE DATA PROCESSING SYSTEM**

BACKGROUND

1. Technical Field

The present invention generally relates to multi-processor data processing systems and in particular to operations on data processing systems configured with multiple independent processing nodes. Still more particularly, the present invention relates to a method and system for completing parallel processing of work items of a single work set distributed across multiple processing units of a multi-node data processing system.

2. Description of the Related Art

Multi-core data processing systems are widely utilized to enable parallel processing of data that can be divided into portions for completion. There are various different topologies of multi-core systems of which non-uniform memory access (NUMA) system topology is one example. To support process scheduling or work scheduling on distributed processing systems such as the NUMA system, separate queues are provided for each processing node because it is assumed that latency (e.g., communication latency, data transfer latency, etc.) between each node is too large or great, according to some metric, to share a common queue for scheduling work. For example, a memory bus (such as a POWER5™ (P5) bus) may operate at a data transfer rate which provides data transfer latency that is too large or great for multiple nodes to share a common queue. Thus, with these types of multi-node processing systems, work processes and associated data must be divided among the separate work queues ahead of work dispatch and execution. Once the execution of work begins in the different processing nodes, a work stealing system/algorithm is then utilized to rebalance the workload in the separate queues. Implementation of these work stealing algorithms injects a large amount of complexity into the scheduler. This complexity can often lead to inefficient run scenarios where work is continuously “balanced” or “re-balanced” between or among two or more nodes.

SUMMARY

Disclosed are a computer system architecture and computer program product for efficient dispatch/completion of a work element within a data processing system having multiple processing nodes exhibiting low access latency to a system memory construct. The method implementation comprises: selecting specific processing nodes or processing units from among the multiple processing nodes to complete execution of a work element that has multiple individual work items that may be independently executed by different ones of the multiple processing nodes and by different ones of the processing units; generating an allocated processor unit (APU) bit mask that identifies at least one of the processing nodes or processing units that has been selected; placing the work element in a first entry of a global command queue (GCQ); associating the APU mask with the work element in the first entry of the GCQ; and responsive to receipt at the GCQ of work requests from each of the multiple processing nodes or the processing units, enabling only the selected specific ones of the processing nodes or the processing units to be able to retrieve work from the work element in the GCQ.

2

The above summary is not intended as a comprehensive description of the claimed subject matter but, rather, is intended to provide a brief overview of some of the functionality associated therewith. Other systems, methods, functionality, features and advantages of the claimed subject matter will be or will become apparent to one with skill in the art upon examination of the following figures and detailed written description.

BRIEF DESCRIPTION OF THE DRAWINGS

The claimed subject matter and equivalents thereof will best be understood by reference to the following detailed description of illustrative embodiments when read in conjunction with the accompanying figures, wherein:

FIG. 1 illustrates an example multi-node, distributed data processing system architecture with multiple nodes each having multiple processor cores and local memory, within which features of one or more of the described embodiments may be implemented;

FIG. 2 illustrates a prior art implementation of a non-uniform memory access (NUMA) system with two processing nodes having local work queues of pre-distributed portions of work, which are periodically re-balanced using a work queue re-balancing algorithm;

FIG. 3 illustrates a first example of a multi-core, multi-node processing system configured with a shared global command queue (GCQ) having an Allocated Processing Units (APU) mask associated with each work entry and processing unit (PU) masks associated with each processor core to enable selective retrieval of work by approved processor cores within the system, in accordance with one embodiment;

FIG. 4 illustrates a work request operation with subsequent retrieval of work items from a work element by processor cores whose PU mask matched the APU mask of the work element, according to one embodiment;

FIGS. 5 (A and B) illustrates a logical comparison of a PU mask of a processor core to an APU mask of a work element within the GCQ to determine which processor cores are allowed to retrieve work items from the work element, according to one embodiment;

FIGS. 6A-6D provides a sequence of block diagrams representing an example global command queue with APU mask and illustrating modifications to count variables within the global command queue as work items are retrieved from a work element for processing by one or more pre-selected processor cores, in accordance with one embodiment;

FIG. 7 is a flow chart of the method by which select processor cores are allocated for dispatching a work element and a corresponding APU mask and count variables are generated and assigned to the work element within the GCQ, in accordance with one embodiment; and

FIG. 8 is a flow chart of the method by which processor cores are verified for allocation of work items from the work element in the GCQ utilizing the APU mask and count variables to track completion of execution of the work items and the work element, according to one embodiment.

DETAILED DESCRIPTION

The illustrative embodiments provide a method, multi-node, multi-core processing system, and computer program product for efficient dispatch/completion of a work element within a data processing system having multiple processing nodes exhibiting low access latency to a system memory construct. The method implementation comprises: selecting specific processing nodes or processing units from among the

multiple processing nodes to complete execution of a work element that has multiple individual work items that may be independently executed by different ones of the multiple processing nodes and by different ones of the processing units; generating an allocated processor unit (APU) bit mask that identifies at least one of the processing nodes or processing units that has been selected; placing the work element in a first entry of a global command queue (GCQ); associating the APU mask with the work element in the first entry of the GCQ; and responsive to receipt at the GCQ of work requests from each of the multiple processing nodes or the processing units, enabling only the selected specific ones of the processing nodes or the processing units to be able to retrieve work from the work element in the GCQ.

In the following detailed description of exemplary embodiments of the invention, specific exemplary embodiments in which the invention may be practiced are described in sufficient detail to enable those skilled in the art to practice the invention, and it is to be understood that other embodiments may be utilized and that logical, architectural, programmatic, mechanical, electrical and other changes may be made without departing from the spirit or scope of the present invention. The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims.

Within the descriptions of the figures, similar elements are provided similar names and reference numerals as those of the previous figure(s). Where a later figure utilizes the element in a different context or with different functionality, the element is provided a different leading numeral representative of the figure number (e.g., 1xx for FIGS. 1 and 2xx for FIG. 2). The specific numerals assigned to the elements are provided solely to aid in the description and not meant to imply any limitations (structural or functional) on the invention.

It is understood that the use of specific component, device and/or parameter names are for example only and not meant to imply any limitations on the invention. The invention may thus be implemented with different nomenclature/terminology utilized to describe the components/devices/parameters herein, without limitation. Each term utilized herein is to be given its broadest interpretation given the context in which that term is utilized. Specifically, the following terms, which are utilized herein, are defined as follows:

- (1) Allocated Processing Units (APU): processor cores that have a bit mask represented (i.e., set to 1) in a global bit mask (APU mask) linked to a specific work element within the global command queue (GCQ). The APUs are the processor cores among all processor cores in the overall system that are allowed to process work from a particular work set within the global command queue.
- (2) Kernel/Work Element: a combination of a function to be executed and the number of times to execute the function over a particular set of input data. Each execution of the function on a unique data is referred to herein as a "work item", and each work element is assumed to comprise multiple work items which can be executed across multiple different processing nodes by multiple different processing units.
- (3) Work Item: a base element of a data set (e.g., a byte, a string, an integer number, a floating point number, a pixel, an array, a data structure, etc.).
- (4) Work Group: a group of work items, which group is referred to herein as a "work element" when placed in a single entry of a global command queue.
- (5) ND Range Size: a number of work items in an index space.

- (6) Work Group Size: a number of work items in a work group, where the number divides into a ND Range Size without a remainder.
- (7) Kernel: a function that processes a work item to produce a result, where the function is implemented by one or more instructions executable by a processor core.
- (8) Kernel Object: a kernel and argument(s) to use when executing the kernel.
- (9) Work Item Remaining (WIR) Counter: a counter that tracks the number of work items within a work element and which is decremented each time one or more work items are removed by a processor core. The amount of the decrement is equal to the number of work items removed by the processor core, and the counter is decremented until the counter value reaches or goes below zero.
- (10) Seen Counter: a counter which tracks the number of approved processing units that have seen/queried the work element for work. The seen counter is initially set to the total number of APUs that are assigned to the work element, and the counter is subsequently decremented each time a different one of the APUs checks the work element for work while the WIR Counter value is at or below zero. A work element is removed from the GCQ whenever the Seen Counter reaches zero.

In the following descriptions, headings or section labels are provided to separate functional descriptions of portions of the invention provided in specific sections. These headings are provided to enable better flow in the presentation of the illustrative embodiments, and are not meant to imply any limitation on the invention or with respect to any of the general functions described within a particular section. Material presented in any one section may be applicable to a next section and vice versa.

A. System Architecture

As provided herein, the invention is applicable to any number of different configurations of such multi-core systems, including ones configured as large scale parallel processing system, or ones configured as non-uniform memory access (NUMA) processing system, or other type of multi-core processing system, including parallel processing systems. The descriptions provided herein assume no specific architecture, other than a multi-core architecture. However, certain features are assumed implemented in multi-node architectures (i.e., system with multiple processing nodes) in which one or more processing nodes exhibit low latency when accessing a global/shared memory structure, while other features are assumed to be implemented in multi-node architectures in which one or more processing nodes exhibit high latency when accessing the global/shared memory structure.

As further described below, implementation of the functional features of the invention is provided within processing nodes and involves use of a combination of hardware, firmware, as well as several software-level constructs. The presented figures illustrate both hardware and software components within an example distributed, multi-node computing environment in which multiple physically separate processing nodes, interconnected via a general system interconnect and/or a network interface device (NID) and/or host fabric interface (HFI) and/or one or more other forms of interconnects (e.g., switches, and the like), provide a distributed data processing system that executes one or more ND work sets via a plurality of processor cores. The illustrative and described embodiments assume that the system architecture may be scaled to a much larger number of processing nodes.

Turning now to the figures, and in particular to FIG. 1, which illustrates an example multi-node, distributed Data

Processing System (DPS) architecture within which the described embodiments may be advantageously implemented. As shown, DPS architecture 100 comprises a local processing system 101, which is coupled to a remote computing device 170 via an interconnect 165, such as a network fabric or switch. Local processing node 101 comprises a plurality of processing nodes of which processing node 110A and processing nodeN 110B are illustrated. Processing node 110A and processing nodeN 110B are coupled to system interconnect 140, which serves as and may interchangeably be referred to as data buses.

System interconnect 140 provides connectivity within processing system 101 to various devices, including but not limited to, memory controllers 150A, 150B and input/output (I/O) controller 156. Network controllers 150A, 150B controls access to respective system memory devices 152A, 152B. Depending on system design, memory controllers 150A, 150B may be assigned to respective processing nodes 110A or 110B, and/or may be physically located on the processor chip or at some other location within the overall system architecture. It is appreciated that the presence of multiple memory controllers 150A, 150B is a design parameter, and that the local processing system 101 may be configured with a single memory controller in an alternate embodiment. In one embodiment, I/O controller 156 provides control over/by one or more I/O devices 158, such as a pointing device, display monitor, and the like.

In addition to system memory devices 152A, 152B of the overall system, each processing node 110A and 110B has an associated node local memory 154A, 154B, which allows for staging of work at the processing node without requiring the processing node to have to retrieve each piece of work individually from the system memory 152A, 152B or other remote (or high latency) location. Each processing node 110A, 110B has a plurality of processor cores, although it is possible for the nodes to be single core nodes. As shown, processor node 110A comprises N+1 cores, numbered core_0 through core_N, where N is an integer greater than zero. Likewise, processor node 110B comprises M+1 cores, numbered core_0 through core_M, where M is an integer greater than zero. So as to simplify the description of the embodiments, both N and M are assumed to be integer three (3), such that processing node 110A has four (4) processor cores 111-114 and processing node 110B has four (4) processor cores 115-119. Where utilized herein, the term processing units are assumed to be synonymous with processor cores. Depending on implementation, the processor cores may be one or more of central processing units (CPUs), graphical processing units (GPUs), synergistic processor unit (SPUs), and other types of processing units.

The use of the integer variables “N” and “M” to cap the number of processor cores in respective processing nodes 110A and 110B merely indicates that each processing node can include a variable number of separate processor cores, with N and M being integers of any size greater than 0 (assuming a multi-core implementation for respective processing nodes). The number of processing nodes deployed in a given system architecture is implementation-dependent and can vary widely. It is appreciated that the embodiments described herein enables scaling to a much larger number of processing nodes and even larger number of processor cores. Within the described embodiments, local processing system 101 is assumed to be the source of the ND Range work set that is generated and processed by the processor cores within DPS architecture 100.

As further illustrated, local processing system 101 also includes storage 120, within which are stored several of the

firmware and software components that enable generation of work and ND Range work sets for execution by one or more of the processing units. Provided within storage 120 are several operating systems (OSes) 122-124, applications 126-127, compiler 125, GCQ scheduler 130, ND Range Work Scheduling (WS) logic 131, and ND range parameters 133 (which may also be stored in system memory 152). In one embodiment, GCQ scheduler 130 is logic that is executed by a processing unit to receive work kernels from a compiler and schedule the work for dispatch to/by one or more of the multiple processor cores within the DPS architecture 100. In one embodiment, the work set that is operated on is retrieved by compiler 125 from one or more of applications 126-127 and/or OSes 122-124, or from some other source (not illustrated herein). Further, local processing system 101 includes global command queue (GCQ) 320, which may be stored in storage 120 or within system memory 152A. Local storage of GCQ 320 enables low latency access to GCQ by processor cores, such as cores 0-N and 0-M, when these processor cores are looking for work to complete. In one embodiment, compiler 125 includes a just in time (JIT) compiler that produces one or more kernels in accordance with a method or process associated with processing work. For example, application 126 can implement a method or process associated with processing work, and in response to an execution of application 126, the JIT compiler can produce one or more kernels to process a data set or work. While generally illustrated as software implemented components, the functionality of the GCQ logic, WS logic, compiler, and other functional logic generally described herein (e.g., LCQ logic) may be implemented as hardware or a combination of hardware and software, in some embodiments.

Local processing system 101 includes or is coupled to a switch or other global interconnect (e.g., interconnect 165) to which multiple other processing nodes may be connected. As illustrated, local processing system 101 is communicatively connected to remote computing device 170 via interconnect 165, which may be a network fabric or switch. Connection to interconnect 165 is provided via network interface controller (NIC) 160, which may be a host fabric interface (HFI). Interconnect 165 may be a single general interconnect to which all nodes connect or may comprise one or more subnets (not shown) connecting a subset of processing nodes to each other and/or to local processing system 101, in one of several alternate configurations.

Remote computing device 170 provides additional computing resources within DPS architecture 100, such as remote processing node 110C and associated local memory 154C. Remote processing node 110C comprises at least one processor core, of which core 171 is provided. Where processing node 110C is a multi-core processing node, processing node 110C may comprise P+1 cores, of which core_P 172 is illustrated. With this multi-core configuration, “P” is an integer greater than or equal to one. For purposes of describing the illustrative embodiments, P is assumed to be one (1), such that processing node 110C has two (2) processing cores. Also, because of the relative distance between remote computing device 170 and specifically processing node 110C from the work staging area within local processing system 101, core_0 171 and core_P 172 are described in some embodiments as cores exhibiting high latency during shared work retrieval and/or balancing, as will become clearly in the descriptions which follow.

The illustrated configuration of multi-node DPS architecture multiple processing nodes is presented herein for illustrative purposes only. The functional features of the embodiments described herein therefore apply to different

configurations of data processing systems. Those skilled in the art will appreciate that DPS architecture **100** of FIG. **1** can include many additional components, such as interconnect bridges, ports for connection to networks or attached devices, and the like (none of which are illustrated herein). Because such additional components are not necessary for an understanding of the present invention, they are not illustrated in FIG. **1** or discussed further herein.

B. Architecture for Low Latency Work Retrieval by Processing Units

Certain of the features of the described embodiments may lend themselves to N Dimensional (ND) Range processing/execution paradigms, such as provided by OpenCL™ and similarly structured applications executed on multi-core data processing systems in which an execution command contains the information necessary to efficiently execute and balance work load of the command, which may include data or task parallel processing requirements. Each command includes information corresponding to the total number of work groups within the command to be executed. The compiler of the processing system receives the command and generates kernels of work, which are scheduled on one or more of the available processor cores of the processing system. With a conventional multi-node processing system, the scheduling of work items associated with each kernel may be completed as entire work groups, utilizing local node queues as illustrated by FIG. **2**.

FIG. **2** thus illustrates a conventional configuration of a multi-node processing system **200** in which work generated by a command is executed by processor cores on different processing nodes **210A**, **210B** by dividing the data set into work groups and arbitrarily and statically assigning the work groups to separate, structured command queues (i.e., pre-set sized queues) that are physically located within each node. As shown, processing node **210A** includes command queue **225A**, while processing node **210B** includes command queue **225B**. Scheduler **215** assigns different portions of the entire work, perhaps based on memory affinity of the addressing scheme of node local memory **254A**, **254B** relative to the data elements within the work, without full consideration for load balancing across the nodes. To complete the dispatch of work and account for dependencies, etc. the nodes have to communicate with each other and may have to move work around the different command queues. There is considerable amount of cross chatter between the nodes to provide coordination with each other as the work is being performed out of their respective local command queues. This configuration leads to work contention over the system interconnect **140** as the processing cores operate at different speeds and often have to rebalance the work load amongst the nodes, leading to substantial amounts of inter-node chatter **230** as the nodes attempt to move work around between the different work queues (**225A**, **225B**).

With reference now to FIG. **3**, there is provided one embodiment of a multi-node processing system architecture **300**, which includes software and hardware constructs to allow for effective and efficient work allocation to and completion by processing nodes and specifically processor cores (on these processing nodes) when the processor cores exhibit low latency communication with system memory of local processing system **300** on which the work is being dispatched. In addition to the hardware structure presented in system architecture (**100**) of FIG. **1**, which have been previously described, FIG. **3** provides a different configuration and functionality of several of the software and firmware compo-

nents introduced in FIG. **1** as well as two new constructs, which together enable the functional features of the described embodiments.

Multi-node processing system architecture **300** is illustrated with three processing nodes, processing node A **310A** and processing node B **310B**, each having four (4) processor cores (nodeA cores **111-114** and nodeB cores **115-118**) and processing node C, with two processor cores (nodeC cores **171-172**). Each processing core within system architecture **300** (namely processing cores **111-114**, processing cores **115-118**, and processing cores **171-172**) is provided a unique bit mask identifier (MB), which is associated/linked/encoded with each processor core and uniquely identifies the specific processor core from other processor cores within the overall system architecture **300**. Thus, for example, core **111** is tagged with first bit mask (MB) **311**, core **112** is tagged with second bit mask **312**, core **114** is tagged with fourth bit mask **314**, core **116** is tagged with sixth bit mask **316**, and so on. For simplicity, the bit masks are referred to herein as processing unit (PU) mask, such that each processing unit/processor core has a unique mask identifier. Example PU masks as well as their functionality are provided in FIG. **4**, which is described below. In one embodiment, each processor node (**310A**, **310B**, **310C**) is also provide a node-level bit mask, nodeA mask (NM) **360A**, nodeB mask **360B**, and nodeC mask **360C** to enable node level scheduling, as described further below.

Turning now to the functional blocks in the upper portion of the figure, the executable/execution code of example application **126** is forwarded to compiler **305**, which processes the instructions/codes to generate commands. These commands provide work divided into work groups, each referred to as kernel **310**. The kernels **310** are forwarded to GCQ scheduler **315**. GCQ scheduler **315** is an enhanced runtime scheduler that dynamically performs the functions of: (1) scheduling (placing) work elements into a global command queue (GCQ) **320**; (2) selective allocating/assigning the work elements to specific selected processor nodes or processing units (cores) from among the entire set of available processor nodes or processing units; generating an APU mask (**330**) corresponding to the specific selected processor cores allocated/assigned the work element; and appending/linking or otherwise associating the APU mask (**330**) to the work element by forwarding the APU mask to GCQ **320**. The above scheduler functions are performed at runtime, and the work items associated with the command are generated at runtime when the work element is placed in GCQ **320** for dispatch. In one embodiment, the selectively allocating of the work elements to selected processor cores involves and/or is based on a scheduling criteria that takes into account workload allocation and work balancing across the system architecture, processing capabilities of the different types of processing units (e.g., CPU, GPU, SPU), and other factors, which may be pre-programmed by the designer of the GCQ scheduler. The actual scheduling of the work elements at the GCQ then involves dispatching work items from the work element in the GCQ to only the selected specific processing nodes or the processing units, wherein the dispatching is responsive to receipt at the GCQ of work requests from each of the multiple processing nodes or the processing units,

The APU mask **320** represents a global bit mask and is a field of 1s and 0s, with the 1s representing the processing units of the node(s) to which the particular work element may be forwarded, and the 0s indicating that a particular processing unit has not been selected to process work for that work element. The actual size of APU mask **320** scales based on the total number of processing nodes and processor cores avail-

able for consideration during allocation of work by GCQ scheduler **315**. In implementations where the division of work is per processing nodes, rather than at the granularity of processor cores, the APU mask represents a node mask, which is a bit mask identifying the specific processor node. Any processor core at that selected processing node may then consume the work items from the work element.

In one embodiment, scheduler assigns a different APU mask to each work element; However, it is contemplated that multiple work elements of a single command (work) may be provided with the same APU mask, leading to an alternate embodiment in which a single APU mask may be utilized for multiple sequential work elements placed in the GCQ. An example GCQ **320** is provided by FIGS. **6A-6D**, which are described below. GCQ **320** may also comprise GCQ logic **335**, which enables selective processor core retrieval of work elements and other aspects of the functionality described hereafter.

It is important to note that the types of work being described herein may be work performed on an N Dimensional (ND) Range of data, which work is passed to the compute nodes, via the GCQ as an ND Range work instruction. In an ND Range instruction, an N dimensional range of work is provided for execution in compute nodes by specific/selected ones of the processor cores, as identified by the APU mask. The value of N in the ND Range and hence the number of dimensions is variable and may be any integer value greater than one (1). According to the described embodiments, the ND Range instructions are executed per kernel invocation, which sets the characteristics of the execution. With the described implementation of GCQ **320**, the APU mask **330** and other execution properties (and/or attributes of the work element) are attached to the work items themselves.

FIG. **6A** illustrates an example GCQ **320**, which comprises multiple queue entries, of which three entries are shown with enqueued work elements, WE₀, WE₁, and WE₂. As provided by FIG. **6A**, in addition to the actual work element, each entry has a number of variables and/or attributes or properties that are associated with (or linked to) the specific work element, and which are stored in one of the pre-established fields of the GCQ entry. The first row of GCQ indicates the nomenclature provided for each field of each of the GCQ entries. Thus each occupied/filled GCQ entry includes the following fields of information: ND Range work element **625**; seen counter **627**; work items remaining (WIR) counter **629**; APU mask **330**; and attributes **631**. For purposes of description, the last/bottom entry is the first entry in GCQ **320** and holds work element zero (WE₀) with its associated properties, attributes, and/or variables. As provided by the WIR count within WIR counter **629** of GCQ **320**, when WE₀ is initially placed within the GCQ entry, WE₀ includes 16 work items within the work set/group generated by compiler **315** (and enqueued within GCQ) that remain to be executed by the selected processing units. Additionally, as indicated by the seen count within seen counter **627**, the 16 work items may be executed by four (4) different selected processing units/cores, which have been identified by scheduler prior to placing the work element in the GCQ. The specific selected processor cores are identified by APU mask **320**, which is a sequence of bits that provide an ordered identification of which processor cores of the N processor cores within the overall system have been assigned to process the 16 work items within WE₀.

In the described embodiments, the start value of the seen counter is presented as the number of allocated processing units to/on which work items of the specific work element may be dispatched or scheduled. However, in an alternate embodiment, the seen count is initiated with a start value

equal to the total number of processing units within the overall system. For example, with system architecture **300** of FIG. **3**, assuming all processing nodes and corresponding processing units have been accounted for, the seen count would be initiated with a value of ten (10). With this alternate embodiment, decrementing of the seen count would be triggered for each processor core that goes to the GCQ for work, regardless of whether the processor core was actually one of the selected processor cores approved for dispatching/scheduling work items from the work element. Either embodiment may be implemented; however, the described embodiments assume the former implementation in which only the approved processor cores are accounted for by the seen count and only those approved processor cores trigger a decrementing by the GCQ logic of the seen count for the particular work element.

Referring now to FIG. **4**, there is illustrated an expanded view of the work items within WE₀ and the associated APU mask **330** which provides access to the work items by selecting ones of the processor cores within example system architecture **300**. As shown, WE₀ comprises 16 individual work items, from work item zero (WI₀) **401** through WI₁₅ **416**. The work items are dispatched to requesting processor cores in a particular dispatch order, which for purposes of illustration, is assumed to be in descending order (WI₁₅ **416**, WI₁₄ **415**, and so on) until the last work item WI₀ is assigned to a processor core. Associated with WE₀ is APU mask **320**, which has a mask value of 10010101. The present example assumes that only processing node A **310A** (FIG. **3**) and processing node B **310B** (FIG. **3**) are represented within APU mask, such that a maximum of eight (8) processor cores from these two processing nodes may be assigned to process WE₀. Processing node C **310C** does not include any approved processors (and the corresponding mask bits are 00 at the end of APU mask).

As shown by the example bit mask in FIG. **4**, with a three node implementation, each of nodes A and B having four processing units/cores and node C having two processing units/cores, APU mask **330** includes 10 bits, one for each of the processing cores, such that if the work element within GCQ is intended to be processed by node A (having cores **0** through **3**), then one of the first four bits in APU mask **330** will be set to 1. Conversely, if the work element is not intended to be processed by any processor cores in node C (having cores **8** and **9**), then all of the last two bits in APU mask **330** will be zero (i.e., the bits are not set to 1).

In the example of FIG. **4**, APU mask **330** provides that four processor cores out of a total of ten (10) possible cores within the system are selected and/or assigned to process work for WE₀. With APU mask **330**, each bit corresponding to each selected processor core is set to one (1), while the bits corresponding to the other processor cores remain zero (0). Thus, with ten processor cores in the overall system, for example, an APU mask value of 1001010100 indicates that processor core **0**, **3**, **5**, and **7** have been selected to process the WE₀. With processing nodes A, B and C of processing system **300** (FIG. **3**), these mask bits correspond to core **0** **111** and core **3** **114** of processing node A **310A** and core **1** **116** and core **3** **118** of processing node B **310B**. For illustrative purposes, these processor cores are also shown as belonging to processing groups, which in one alternate embodiment enables work to be dispatched/grabbed at a processing group level (or node level) rather than the independent grabbing of work by each processor core. It should be noted, that in an alternate embodiment, a node-level allocation of work items can be implemented by which all processor cores in a particular node share the same node mask and the APU mask only tracks selected processor nodes. With the above example, of the three avail-

able nodes within the overall system, node A and node B are the selected processor nodes, and the value of the APU mask would be 110, with each bit representing a separate node. The efficiencies of performing the comparisons of the APU mask against the node masks are clear from the significant reduction in the number of bits that have to be tracked by GCQ logic 335.

FIG. 4 also illustrates part of the process by which the various processor cores attempt to “grab” work from GCQ 320, and specifically WE_0. When a processor core needs work, the core issues a work request, such as work request 440/441 to GCQ 320. The work request (440/441) is received and processed by GCQ logic 335 (FIG. 3). The work requests 440/441 comprise PU mask of the requesting processor core, as well as a chunk size parameter, indicating the amount of work (e.g., number of work items) the processor core would like to grab with the request. GCQ logic 335 and/or work request logic of the requesting processor core compares the PU mask of the processor core with the APU mask to determine whether the processor core may grab work items from WE_0. FIG. 5 provides an illustration of one method of completing this comparison. As provided by FIG. 5, each processor core has an associated PU mask. When the processor core issues a request to the GCQ for work, the GCQ logic 335 performs a logical AND of the PU mask of the processor core with the APU mask 330 to generate a result (i.e., a 1 or 0), where the result (if a 1) indicates that the processor core is approved to grab work from the corresponding work element. According to the example of FIGS. 4 and 5, the AND operation with first operand of APU mask 330 (with value of 10010101) with second operands of different PU masks (311-318) of the eight processor cores from the example system yields a “1” output/result for core0 111 and core3 114 of processing nodeA 310A and core1 116 and core3 118 of processing nodeB 310B.

Thus, as provided by FIG. 4, these approved/assigned processor cores are then able to grab work items in the work item dispatch order for execution by the respective processor core. For illustrative purposes, one or more processor core(s) of processing nodeA 310A grab a first set of work items, while one or more processor core(s) of processing nodeB 310B grabs the next sequential set of work items. However, the order in which processing cores grab work items is not dispositive, as long as the work items are grabbed in a known sequence, which is maintained by the scheduler and/or compiler, to enable re-creation of the output.

In the above described work scheduling and dispatch method and system and according to the described embodiments, execution units process the commands from a single GCQ comprising multiple queue entries. Within the GCQ, each entry holds a work element (or work group) comprising multiple work items that are dispatched for execution by specific processing units identified by the APU mask 330 assigned/linked to the work element. As provided by the described embodiments, the execution/processing units are processing cores, which may be located across a distributed network. However, while processor cores are employed throughout the described and illustrative embodiments as examples of execution units, in other embodiments, the execution units may be any device that executed a command, including, but not limited to, processing codes, CPUs, threads, and even complete computing systems. Once a command is available on GCQ 335, each processor core allowed to participate in the execution of the command (as identified by the APU mask) atomically decrements the command’s work items remaining (WIR) counter by a work reservation size (or chunk size) associated with the processor core. The

processor core reserves a corresponding number of work items within a work group range for processing. Once the processor core has executed the requested work items, the processor core attempts to reserve more work items. This reservation and dispatch of work items in a given reservation size by multiple processor cores continues until the total number of work items within one entry of the command queue(s) have been processed.

With the above functionality of assigning work to processing cores utilizing a GCQ, APU masks and PU masks, the process for scheduling and completing all work involves one or more methods that involve generation and tracking of various parameters to enable efficient work scheduling, dispatch and execution. FIGS. 7 and 8 provide flow charts of example methods for completing the above processes. The methods of FIGS. 7 and 8 are best understood when read in the context of the illustrations of FIGS. 3-5 described above, as well as with reference to FIGS. 6A-6D, which are concurrently described. It is appreciated that the methods may be recursive in nature; However, to simplify the descriptions, only a single iteration of each method is illustrated and described. Various portions of the method(s) are performed by processing unit executing instruction code of the executing task/thread, while other portions may be performed by firmware/logic associated with or executing on the processor. The method(s) is/are generally described from the perspective of processor operations, although it is understood that alternate embodiments of the method may be performed by GCQ logic 335 and GCQ scheduler (logic) 315, for example.

The flowcharts and block diagrams in the figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowcharts or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

The method of FIG. 7 begins at block 702 and proceeds to block 704 at which the compiler generates work elements (kernels or commands) from instruction code provided by a source program. The compiler 305 forwards the commands/kernels 310 to the GCQ scheduler 315. During runtime, the scheduler (logic) 315 selects a set of one or more of the available processing units (processor cores) from a global set of available processing units to execute the command/kernel 310, as shown at block 706. GCQ scheduler 315 also generates the APU mask corresponding to the selected set of available processing units, at block 708. As shown at block 710, GCQ scheduler 315 also assigns (a) a seen count value based on the number of selected processing units and (b) a work item remaining value based on the number of work items associated with the work element. In an alternate embodiment, the WIR count may be generated and updated by GCQ logic 335, perhaps based on a detection of the actual work items scheduled within work element. Finally, as proved at

block 712, GCQ scheduler 315 forwards the APU mask 330, seen count, and WIR value to GCQ 320 when an entry becomes available in GCQ 320. The process then ends at block 714.

At the termination of the above processes, GCQ 320 comprises an entry populated by the work element and associated properties, parameters and attributes. Referring again to FIG. 6A, WE_0 is populated in the first entry of GCQ 320 and is ready for dispatch of its 16 work items by the four approved processing units identified by the APU mask.

Turning now to FIG. 8, which depicts the method by which the work items of WE_0 are dispatched by the processor cores of the system. The method begins at initiation block 802. At block 804, first processing unit (e.g., processor core_0 311) issues a request to the GCQ to check the availability of work. The request provided by first compute unit comprises its PU mask and the chunk size, which are forwarded to GCQ logic 335. At block 806, the processing unit (or the GCQ logic) compares the PU mask with the APU mask of the first work element (and subsequent work elements, if unsuccessful in scheduling work), and determines at decision block 808 whether the comparison of the PU mask with the APU mask indicates that the PU is approved/authorized/pre-selected to grab work items from that work element. According to one embodiment, and as illustrated by FIG. 5, the process by which a processing unit confirms that it is one of the selected processing units to execute a particular work element involves performing a logical AND of the APU mask of the work element and the PU mask of the processor core. If the processing unit is not approved/authorized/pre-selected to execute work items of the work element, GCQ logic 335 forwards the request to a next work element in the GCQ, as shown at block 710. Thus, the request from the processor core moves on up the GCQ to the next entry, and the same check is performed against the APU mask of the next work element with the PU mask in the request. In this manner, a processor core only removes work from a work element when that processor core has been pre-selected by the scheduler to execute work of that work element.

As further illustrated by FIGS. 6D, the processing units that are not approved for processing WE_0 may also concurrently schedule work from a different work element at which the processing unit is approved. For example, processing unit 2 of processing node A 310A may schedule work from WE_1 within GCQ 320, after first being rejected at WE_0. When work items are scheduled on processor core 1, the WIR counter and the seen counter of WE_1 are decremented. As shown, WIR counter of WE_1 is decremented by 10 as ten work items are scheduled on processing unit 2, and the seen count of WE_1 is also decremented by one (1).

When the processor core has been confirmed as an approved/authorized/selected processing unit, GCQ logic 335 then checks the value of the WIR counter, and determines at block 812, whether the WIR count is greater than zero (0). When the WIR count is not greater than zero, GCQ logic 335 decrements the seen count value by one (1), as provided at block 718. When the WIR count is greater than zero (0), GCQ logic 335 allows the processor core request to pull/grab an amount of work items from the work element up to the chunk size identified in the request, as indicated at block 814. GCQ logic 335 monitors the removal/dispatch of the work items from the queue entry and decrements the WIR counter by the number of work items removed, as shown at block 816. The process then proceeds to block 817, which issues a new request when execution of the command on the previously retrieved work items have completed. The same processor core may thus continually issue requests to the GCQ and

retrieve work from the same work element until all work of that work element has been assigned.

It is important to note that the number of work items remaining to be dispatched for execution may be less than the chunk size requested by the processor core. In this scenario, all of the remaining work items are provided to the requesting processor core, and the WIR counter is reduced to zero (0). Also, in an alternate embodiment, the chunk size is pre-set within the GCQ, such that either (a) all processor cores are allowed to take the same chunk size of work or (b) larger chunk sizes can be requested but are scaled as whole number multiples of the pre-set chunk size. In the presently described embodiment, the chunk size is parameter set by the processor core (or processing node) and encoded within the request.

The value of the WIR counter is utilized to determine when more work items are available for grabbing, and is decremented each time work items are allocated to a requesting, approved processor. As indicated by the sequence of FIGS. 6A-6D, as each approved processor removes its chunk size of work from the work element, the GCQ logic decrements the WIR counter by the number of individual work items removed (e.g., the chunk size). For illustrative purposed, a singular chunk size is assumed of four (4) work items. Thus (referencing WE_0 of FIG. 4), a request issued by a first processor unit (e.g., request 440 from processor core 0 311) removes work items 15 through 12 from WE_0. GCQ logic 335 decrements WIR counter from 16 to 12 (FIG. 6B). Similarly, when second processor unit request 442 from processor core 1 316 removes work items 11 through 8 from WE_0, GCQ logic 335 decrements WIR counter to 8, and so one, until WIR counter equals or goes below zero (0) (FIG. 6C).

The above method processes describe the GCQ scheduler and/or GCQ logic queuing a work element within the GCQ and the processing elements subsequently removing work from the work element for execution. These processes include/involve GCQ logic 335 (or other logic) updating/modifying certain parameters/attributes of/associated with the work element to enable sequenced retrieval of different work items by each approved processing unit that submits a request for work. An additional mechanism is provided to also track when dispatching of all work from the work element has completed, in order to allow the work element to be removed/discarded from the GCQ.

Work elements are released from the command queue when all the cores within the overall system or all the cores selected to dispatch the work element (depending on implementation) have seen the work element and do not need to process the work element further. Tracking of the number of processing units that have seen the work element is provided by an atomic decrement of a seen counter that begins with the total number of processor cores, and is decremented each time a different one of the processing units "looks at" (i.e., attempts to schedule work from) the work element. When the last processor core has seen the work element, the seen counter is decremented to zero and the last processor core then notifies GCQ scheduler 315 or GCQ logic 335 that execution of the work element has completed.

Returning now to block 812 of FIG. 8, when the WIR counter value of a work element is zero (or less than zero), indicating that no more work items remain to be dispatched, and a request from an approved processing unit is received, GCQ logic 335 decrements the seen counter by one (1), as provided at block 818. At block 820 of the flow chart, a decision is made whether the seen count equals zero (0). If not, GCQ logic 335 monitors for receipt of subsequent requests from each other approved processing unit, and decrements seen count for each request received from a different,

approved processing unit. Once all the approved processing units have submitted requests that are rejected because there is no more work to dispatch, the seen count is finally decremented to zero. The decrementing of seen count for WE_0 is illustrated by FIGS. 6C-6D. Once seen count becomes zero, the work element is removed from or overwritten in the GCQ 320, as shown at block 822. The process then ends at block 824.

Features of the provided embodiments can advantageously be implemented in multi-node processing systems which exhibit low data transfer latencies and low cross communication latencies between the processor cores and a memory location of the local node. Given that the processing nodes are retrieving work from the central/common GCQ, the latency to the GCQ for the different nodes is relative short, such that the overall work retrieval process incurs an acceptable latency penalty in overall processing time for the processing cores that are grabbing work from the GCQ 320.

With the implementation of the single command queue, the need to synchronize across multiple command queues of multiple nodes is substantially eliminated. The described embodiments are adapted for implementation on multi-processor compute systems (or system architecture) that exhibit small penalty (latency) when synchronizing across compute nodes. For example, with NUMA nodes that are connected via a "local" fabric, such as within a single physical device (e.g., local processing system 301), there is a relatively small latency penalty when synchronizing distributed work of a single kernel (work element) across the multiple nodes. This further enables the completion of N Dimensional work and data intensive operations to/on NUMA systems with low cross node latency.

Also, with the described embodiments, there is no longer a requirement for multiple individual command queues that are each tied to a separate node in order to support NUMA affinity. With the single GCQ implementation, each work element within the single GCQ includes a mask field comprising N bits, with specific bits set to identify which execution nodes and/or which processing unit(s) (or processor cores) are linked to the work element for completion thereof. The N bit mask provides a unique identification of each separate node and/or processor core so that the same single command queue can be utilized across multiple nodes, while allowing the nodes/cores to correctly grab work assigned to the particular node/core from the GCQ. Because the present description of one embodiment includes specific reference to a NUMA system architecture, the N bit mask may also be referred to within this embodiment as a NUMA mask, and each executable element (e.g., core_0 311) is assigned a NUMA mask that is unique for the particular node or the particular processing unit (depending on the level of granularity desired). When a new command is placed in an entry of the single GCQ, the new command (work element) is marked/tagged/assigned/affiliated or otherwise associated with a specific NUMA affinity via the NUMA mask.

The above embodiments are applicable to systems that have NUMA regions close enough so they can share the GCQ efficiently. This is especially true for workloads where a common function is being executed multiple, independent times. Each work element in the GCQ contains an APU mask, work remaining count (in the case of multiple iterations of a process), and a seen count. With these functional variables and system configuration, a work element within the GCQ is scheduled as follows: (a) the processing unit checks if the processing unit's NUMA bit is set in the next work element's APU mask; (b) if the processing unit's NUMA bit is not set in the APU mask, the processing unit decrements the seen count,

indicating that the unit has finished with the work element, and the processing unit advances to the next element; (c) If the processing unit's NUMA bit is set, the processing unit decrements the work item remaining (WIR) count by a tunable "chunk size", based on the actual amount of work retrieved from the work element by the processing unit; (d) If the result of the decrement is a value indicating that there is no more work (e.g., 0 remaining work items), the processing unit decrements the seen count, which indicates that the processing unit has finished with the element. The processing unit then advances to the next work element; (e) if the processing unit is the last processor core to see the element (i.e., seen count is equal to zero), indicating all the processor cores have seen the work item at least one, the processing unit marks the location in the GCQ as "free."

At the processing unit (processor core), the processing unit executes each work item in that retrieved chunk until the execution of the chunk of work is completed. The processing unit then issues a next request to the queue to get more work.

Certain of the features of the described embodiments may lend themselves to N Dimensional (ND) Range processing/execution paradigms, such as provided by OpenCL™ and similarly structured applications executed on multi-core data processing systems in which an execution command contains the information necessary to efficiently execute and balance work load of the command, which may include data or task parallel processing requirements. Each command includes information corresponding to the total number of work groups within the command to be executed, the number of work groups remaining to be computed, the number of work groups already processed, the number of work groups to process at a time (reservation size). Other information provided by the commands may include the number of execution threads to process in the command and the number of execution threads that have seen the command.

Within this ND Range execution paradigm, and according to the presently described embodiments, execution units process the commands from a single global command queue comprising multiple queue entries. Within the GCQ, each entry holds a work element comprising the command and multiple work items that are dispatched for execution by specific processing units identified by an APU mask assigned/linked to the work element. As provided by the described embodiments, the execution units are processing cores, which may be located across a distributed network. However, while processor cores are employed throughout the described and illustrative embodiments as examples of execution units, in other embodiments, the execution units may be any device that executed a command, including, but not limited to, processing codes, CPUs, threads, and even complete computing systems. Once a command is available on the global command queue, each processor core participating in the execution of the command (as identified by the APU mask) atomically decrements the command's work items remaining (WIR) counter by a work group reservation size (or chunk size) associated with the processor core. The processor core reserves a corresponding number of work items within a work group range for processing. Once the processor core has executed the requested work items, the processor core attempts to reserve more work items. This reservation and dispatch of work items in a given reservation size by multiple processor cores continues until the total number of work items within one entry of the command queue(s) have been processed.

In each of the flow charts above, one or more of the methods may be embodied in a computer readable medium containing computer readable code such that a series of steps are

performed when the computer readable code is executed on a computing device. In some implementations, certain steps of the methods are combined, performed simultaneously or in a different order, or perhaps omitted, without deviating from the spirit and scope of the invention. Thus, while the method steps are described and illustrated in a particular sequence, use of a specific sequence of steps is not meant to imply any limitations on the invention. Changes may be made with regards to the sequence of steps without departing from the spirit or scope of the present invention. Use of a particular sequence is therefore, not to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims.

As provided herein, the embodiments describe a method, computer system, and computer program product for efficient dispatch/completion of a work element within a data processing system having multiple processing nodes exhibiting low access latency to a system memory construct. The method implementation comprises: selecting specific processing nodes or processing units from among the multiple processing nodes to complete execution of a work element that has multiple individual work items that may be independently executed by different ones of the multiple processing nodes and by different ones of the processing units; generating an allocated processor unit (APU) bit mask that identifies at least one of the processing nodes or processing units that has been selected; placing the work element in a first entry of a global command queue (GCQ); associating the APU mask with the work element in the first entry of the GCQ; and responsive to receipt at the GCQ of work requests from each of the multiple processing nodes or the processing units, enabling only the selected specific ones of the processing nodes or the processing units to be able to retrieve work from the work element in the GCQ.

As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

A computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, R.F, etc., or any suitable combination of the foregoing. Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

Aspects of the present invention are described below with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks. The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

As will be further appreciated, the processes in embodiments of the present invention may be implemented using any combination of software, firmware or hardware. As a preparatory step to practicing the invention in software, the pro-

19

programming code (whether software or firmware) will typically be stored in one or more machine readable storage mediums such as fixed (hard) drives, diskettes, optical disks, magnetic tape, semiconductor memories such as ROMs, PROMs, etc., thereby making an article of manufacture in accordance with the invention. The article of manufacture containing the programming code is used by either executing the code directly from the storage device, by copying the code from the storage device into another storage device such as a hard disk, RAM, etc., or by transmitting the code for remote execution using transmission type media such as digital and analog communication links. The methods of the invention may be practiced by combining one or more machine-readable storage devices containing the code according to the present invention with appropriate processing hardware to execute the code contained therein. An apparatus for practicing the invention could be one or more processing devices and storage systems containing or having network access to program(s) coded in accordance with the invention.

Thus, it is important that while an illustrative embodiment of the present invention is described in the context of a fully functional computer (server) system with installed (or executed) software, those skilled in the art will appreciate that the software aspects of an illustrative embodiment of the present invention are capable of being distributed as a program product in a variety of forms, and that an illustrative embodiment of the present invention applies equally regardless of the particular type of media used to actually carry out the distribution.

While the invention has been described with reference to exemplary embodiments, it will be understood by those skilled in the art that various changes may be made and equivalents may be substituted for elements thereof without departing from the scope of the invention. In addition, many modifications may be made to adapt a particular system, device or component thereof to the teachings of the invention without departing from the essential scope thereof. Therefore, it is intended that the invention not be limited to the particular embodiments disclosed for carrying out this invention, but that the invention will include all embodiments falling within the scope of the appended claims. Moreover, the use of the terms first, second, etc. do not denote any order or importance, but rather the terms first, second, etc. are used to distinguish one element from another.

The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. As used herein, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising," when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the present invention has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the invention. The embodiment was chosen and described in order to best explain the principles of the inven-

20

tion and the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.

What is claimed is:

1. A data processing system comprising:

a plurality of processing nodes, each processing node having at least one processing unit;

a system memory construct coupled to the plurality of processing nodes via an interconnect;

a global command queue (GCQ) maintained within the system memory construct during data processing operations utilizing multiple processing resources; and

a scheduler performing the functions of:

selectively allocating a work element having multiple individual work items to specific processing nodes or processing units from among the plurality of processing nodes in order to complete execution of the work element, wherein the multiple individual work items may be independently executed by different ones of the plurality of processing nodes and by different ones of the processing units;

generating an allocated processor unit (APU) bit mask that identifies at least one of the specific processing nodes or processing units that has been selectively allocated the work element, wherein the APU bit mask further comprises a bit for each processing node and processing unit from among the plurality of processing nodes;

in response to generating the APU bit mask:

placing the work element in a first entry of the GCQ; and

associating the APU bit mask with the work element in the first entry of the GCQ; and

in response to the GCQ receiving work requests from each of the plurality of processing nodes or the processing units, dispatching the multiple individual work items from the work element in the GCQ to only the selected specific processing nodes or processing units.

2. The data processing system of claim 1, wherein said dispatching the multiple individual work items from the work element in the GCQ to only the selected specific processing nodes or the processing units, further comprises:

a logic of the GCQ comparing a processing unit (PU) mask of a first processor node or a first processing unit with the APU bit mask;

in response to a comparison of the PU mask with the APU bit mask generating a first result that indicates that the processing node or first processing unit is a selected execution unit for the work element, providing up to a pre-set amount of the multiple individual work items to the first processor node or first processing unit in response to the work request; and

in response to a comparison of the PU mask with the APU bit mask generating a second result that indicates that the first processor node or first processing unit is not a pre-selected execution unit for the work element, forwarding the work request to a next entry of the GCQ holding a next work element with a next APU bit mask;

wherein the pre-set amount is a chunk size desired by the first processor node or first processing unit, which chunk size is identified in the work request received.

3. The data processing system of claim 2, wherein the GCQ logic performs the functions of:

providing, within a work items remaining (WIR) counter associated with the first entry of the GCQ, an initial WIR

21

counter value corresponding to a number of individual work items of the work element that is placed within the first entry of the GCQ; and
 in response to the comparison of the PU mask with the APU bit mask generating the first result, and the pre-set amount of the multiple individual work items have been dispatched to the first processor node or first processing unit, dynamically decrementing a current WIR counter value of the WIR counter by an amount equal to the number of work items provided to the first processor node or first processing unit.

4. The data processing system of claim 3, wherein the scheduler and the GCQ collaborate to provide the functions of:

providing, within a seen counter associated with the first entry of the GCQ, an initial seen counter value corresponding to a total number of the selected specific processing nodes or processing units, and which tracks detection of a work request at the GCQ from each of the specific processing nodes or processing units that is checked against the first entry of the GCQ;

monitoring the WIR counter to determine when a current WIR counter value is equal to or less than zero (0); and in response to the current WIR counter value being equal to or less than zero (0), and a second work request is received from one of the pre-selected execution units, automatically decrementing a current seen counter value by one (1).

5. The data processing system of claim 4, wherein the GCQ logic provides the functions of:

detecting when the current seen counter value is equal to zero (0);

in response to the current seen counter value is equal to zero (0), marking the first entry as available for re-allocation to a next work element; and

signaling the runtime scheduler with an availability of the first entry for allocation to the next work element.

6. The data processing system of claim 2, wherein the comparing of the PU mask with the APU bit mask further comprises performing a logical AND of the PU masks and the APU bit mask, wherein the logical AND yields a first positive result when the PU mask corresponds to one of the allocated specific processor node or processing unit and a second negative result when the PU mask does not correspond to one of the allocated specific processor nodes or processing units.

7. The data processing system of claim 2, wherein the PU mask is a processor node mask and all processing units operating within a single processing node shares the same PU mask and are able to request and receive work items for scheduling at the specific processing unit.

8. The data processing system of claim 1, further comprising:

a compiler and work scheduling (WS) logic that performs the functions of:

receiving one or more instruction codes for generating a command for execution over of an N Dimensional (ND) Range of data;

in response to determining the number of dimensions in the ND Range is greater than one (1), automatically re-configuring/converting a ND Range input data set into a one (1) dimension array comprising N work items that are sequentially ordered to represent a single one dimensional work set;

placing the single work set within one or more GCQ entries as a 1D array of work items that are individually selected in sequential order for execution by the pre-selected execution units;

22

storing an ordering sequence utilized to re-configure/convert the ND Range into the 1D Range array;
 storing a total number of work items within the 1D array of work items;

storing a dimensionality of the ND Range; and re-configuring an output from execution of the 1D array of work items into an ND range output utilizing the stored dimensionality, total number of work items, and ordering sequence.

9. The data processing system of claim 1, wherein the APU bit mask is associated with a plurality of work elements in the GCQ.

10. The data processing system of claim 1, wherein each of the work requests further comprises a chunk parameter that indicates a number of one or more work items a corresponding processor is requesting.

11. The data processing system of claim 1, wherein the specific processing nodes or processing units further comprise at least one graphical processing unit (GPU) and at least one synergistic processor unit (SPUs).

12. An article of manufacture embodied as a computer program product comprising:

a non-transitory computer readable storage medium; and program code on the non-transitory computer readable storage medium that when executed by a first processing node of a data processing system having a plurality of processing nodes, each with at least one processing unit, provides a method for efficient dispatch/completion of a work element having multiple individual work items that are independently executable, the method comprising:

selectively allocating the work element to specific processing nodes or processing units from among the plurality of processing nodes in order to complete execution of the work element, wherein the multiple individual work items may be independently executed by different ones of the plurality of processing nodes and by different ones of the processing units;

generating an allocated processor unit (APU) bit mask that identifies at least one of the specific processing nodes or processing units that has been selectively allocated the work element, wherein the APU bit mask further comprises a bit for each processing node and processing unit from among the plurality of processing nodes;

in response to generating the APU bit mask:
 placing the work element in a first entry of the GCQ;
 and

associating the APU bit mask with the work element in the first entry of the GCQ; and

in response to the GCQ receiving work requests from each of the plurality of processing nodes or the processing units, dispatching the multiple individual work items from the work element in the GCQ to only the selected specific processing nodes or processing units.

13. The article of manufacture of claim 12, wherein said dispatching work items from the work element in the GCQ to only the selected specific processing nodes or the processing units, further comprises:

comparing a processing unit (PU) mask of a first processor node or a first processing unit with the APU bit mask;

in response to a comparison of the PU mask with the APU bit mask generating a first result that indicates that the processing node or first processing unit is a selected execution unit for the work element, providing up to a pre-set amount of the multiple individual work items to

23

the first processor node or first processing unit in response to the work request; and
 in response to a comparison of the PU mask with the APU bit mask generating a second result that indicates that the first processor node or first processing unit is not a pre-selected execution unit for the work element, forwarding the work request to a next entry of the GCQ holding a next work element with a next APU bit mask;
 wherein the pre-set amount is a chunk size desired by the first processor node or first processing unit, which chunk size is identified in the work request received.

14. The article of manufacture of claim 13, wherein the method further comprises:

- providing, within a work items remaining (WIR) counter associated with the first entry of the GCQ, an initial WIR counter value corresponding to a number of individual work items of the work element that is placed within the first entry of the GCQ;
- in response to the comparison of the PU mask with the APU bit mask generating the first result, and the pre-set amount of the multiple individual work items have been dispatched to the first processor node or first processing unit, dynamically decrementing a current WIR counter value of the WIR counter by an amount equal to the number of work items provided to the first processor node or first processing unit;
- providing, within a seen counter associated with the first entry of the GCQ, an initial seen counter value corresponding to a total number of the selected specific processing nodes or processing units, and which tracks detection of a work request at the GCQ from each of the specific processing nodes or processing units that is checked against the first entry of the GCQ;
- monitoring the WIR counter to determine when a current WIR counter value is equal to or less than zero (0);
- in response to the current WIR counter value being equal to or less than zero (0), and a second work request is received from one of the pre-selected execution units, automatically decrementing a current seen counter value by one (1);
- detecting when the current seen counter value is equal to zero (0);
- in response to the current seen counter value is equal to zero (0), marking the first entry as available for re-allocation to a next work element; and
- signaling the runtime scheduler with an availability of the first entry for allocation to the next work element.

24

15. The article of manufacture of claim 13, wherein: the comparing of the PU mask with the APU bit mask further comprises performing a logical AND of the PU masks and the APU bit mask, wherein the logical AND yields a first positive result when the PU mask corresponds to one of the allocated specific processor node or processing unit and a second negative result when the PU mask does not correspond to one of the allocated specific processor nodes or processing units; and the PU mask is a processor node mask, all processing units operating within a single processing node shares the same PU mask and are able to request and receive work items for scheduling at the specific processing unit.

16. The article of manufacture of claim 12, the method further comprising:

- receiving one or more instruction code for generating a command for execution over of an N Dimensional (ND) Range of data;
- in response to determining the number of dimensions in the ND range is greater than one (1), automatically re-configuring/convert a ND Range input data set into a one (1) dimension array comprising N work items that are sequentially ordered to represent a single one dimensional work set;
- placing the single work set within one or more GCQ entries as a 1D array of work items that are individually selected in sequential order for execution by the pre-selected execution units;
- storing an ordering sequence utilized to re-configure/convert the ND Range into the 1D Range array;
- storing a total number of work items within the 1D array of work items;
- storing a dimensionality of the ND Range; and
- re-configuring an output from execution of the 1D array of work items into an ND range output utilizing the stored dimensionality, total number of work items, and ordering sequence.

17. The article of manufacture of claim 12, wherein the APU bit mask is associated with a plurality of work elements in the GCQ.

18. The article of manufacture of claim 12, wherein each of the work requests further comprises a chunk parameter that indicates a number of one or more work items a corresponding processor is requesting.

19. The article of manufacture of claim 12, wherein the specific processing nodes or processing units further comprise at least one graphical processing unit (GPU) and at least one synergistic processor unit (SPUs).

* * * * *